



U.S. DEPARTMENT OF
ENERGY

Office of
Science

SC-HEP Allocations at NERSC

Eric Church, HEP Detailee (PNNL)

NERSC Users Group meeting

12-Oct-2022

High Energy Physics

DOE High Energy Physics Mission is to understand how the universe works at its most fundamental levels by:

- **Discovering** the most elementary constituents of matter and energy
- **Probing** the interactions between them
- **Exploring** the basic nature of space and time

► **HEP develops and supports a specific portfolio of projects:**

- Makes **significant, coherent contributions to facilities/experiments** selected for the program, including project management
- Supports **R&D that will advance the state-of-the-art in particle accelerators and detectors** that will lead to new, more capable facilities
- Supports **R&D to enable new and transformative capabilities** in AI/ML, QIS, and cross-cutting technology areas
- Supports vibrant theory program to **provide the vision and extend our knowledge** of particles, forces, space-time, and the universe

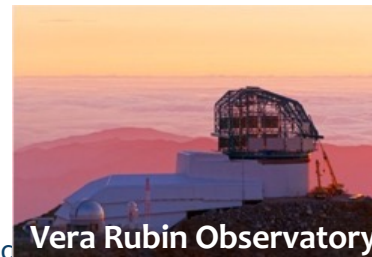
► **DOE-HEP supports ~85% of U.S. particle physics (in \$), including ~all national laboratories**



CMS at CERN



Fermilab Muon g-2



Vera Rubin Observatory

August 2021

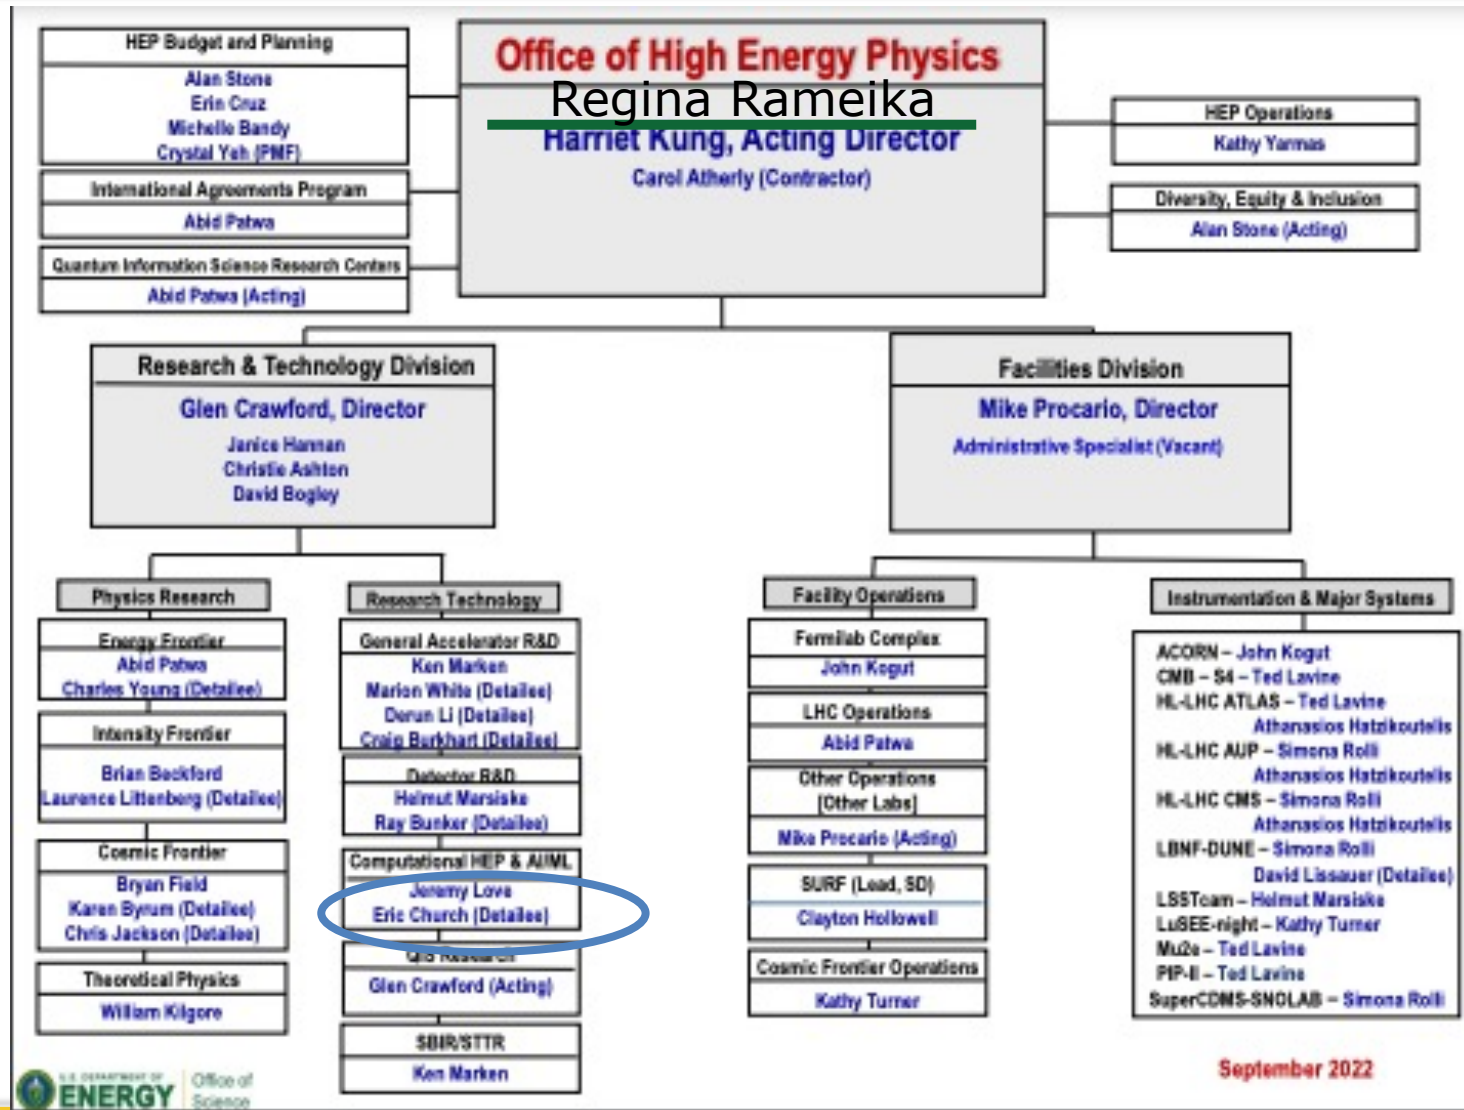
FESAC Meeting



U.S. DEPARTMENT OF
ENERGY

Office of
Science

HEP Org chart



HEP programs

▶ The High Energy Physics Programs

- ▶ Intensity Frontier
- ▶ Energy Frontier
- ▶ Cosmic Frontier
- ▶ Theory

Core Research
Programs

- ▶ ... Accelerators, Advanced Detectors, SBIR, QIS ...
- ▶ Experiment Operations
- ▶ ... and **Computational Physics** – Us! -- manages compute and AI/ML projects ... and things like HEP's presence at NERSC
 - ▶ Jeremy Love, Program Director
 - ▶ Eric Church, detailee



Talk Outline

- ▶ New Allocation Year (AY) applications
 - ▶ Via ERCAP
 - ▶ Will focus on AY22-23
- ▶ Intra-AY Rebalancing
 - ▶ Via IRIS
 - ▶ Mainly AY22 experience
- ▶ Out-Year Planning
 - ▶ For our current large consumers
 - ▶ Doesn't really make accommodation for new large consumers, who we know are coming ...



Allocation Year ERCAP Proposals

- ▶ In October the ERCAP deadline closes and ~80 applications for AY22-23 are filed which ask for an HEP allocation
- ▶ Jeremy and I filter them for relevance to HEP
- ▶ Some things which are not obviously in the program mission may be filtered at this step
- ▶ e.g., esoteric astrophysics studies
- ▶ Gravitational Waves work
- ▶ Mis-targeted things
 - ▶ FES, NP proposals, for example



We are dominated by giants

- ▶ Our top 10 projects use 85% of our allocation.
- ▶ Our top 3 would happily gobble up the entire HEP allocation



AY23's proposals

< ≡ ERCAP DOE Office Program - HEP - High Energy Physics ...

Description	2023 - HEP - High Energy Physics	Has Reserve	<input checked="" type="checkbox"/>
Allocation Year	2023	CPU Initial Reserve	2,638,500
Program Id	15646	CPU Hours Requested	3,196,022
Program Name	High Energy Physics	CPU Hours Awarded	0
Program Label	HEP	CPU Remaining Reserve	2,638,500
Program - Subprogram	HEP - High Energy Physics	CPU % Distributed	0
Subprogram Id	38878	Approval Group	DOE-SC HEP Mgrs ⓘ
Subprogram Name	High Energy Physics	Restrict to Admin	<input type="checkbox"/>

Note that we're being asked for more CPU node hours than HEP has.

This form continues on below what's shown us. Our GPU requests sum to almost what's available for us this year to allocate out.



CPU node hours

- ▶ Those numbers for CPU node hours for us to allocate out (requested) is down from AY2022:
 - ▶ AY23: 2.6(**3.2**) E6
 - ▶ AY22: 4.11(5.38) E6.
- ▶ This is the effect of Cori going away
 - ▶ For which we have tried to prepare the HEP community
 - ▶ The HEP CPU request is down not because the need for CPU time is lower but because we have gone through the planning to manage the pain of Cori's decommissioning. NERSC could not meet the CPU needs of HEP in AY22 and it is really falling short in AY23.



AY22 requests

Just a snapshot of the ~80 proposals



<input type="checkbox"/>	i	Physics : High Energy Physics (Experimen...	ERCAP0021161	2022	Nugent, Peter (nugent)	Berkeley Lab	CCE Training	m2845	m2845	The HEP-CCE is a cross-cutting initiativ...	DOE Mission Science	Provisioned
<input type="checkbox"/>	i	Physics : Cosmology	ERCAP0021329	2022	Lukic, Zarija (zarija)	Berkeley Lab	Large-Scale Structure of the Universe	CosmoBears	m3058	This is a repository used by students, p...	DOE Mission Science	Provisioned
<input type="checkbox"/>	i	Physics : Cosmology	ERCAP0020008	2022	Diegel, Seth (diegel)	SLAC	LSST Dark Energy Science Collaboration (...)	LSST_DESC	m1727	The Vera C. Rubin Observatory's Legacy S...	DOE Mission Science	Provisioned
<input type="checkbox"/>	i	Physics : Accelerator Science	ERCAP0020443	2022	Ng, Cho-Kuen (cho)	SLAC	Advanced Modeling for Particle Accelerators	AMPA	m349	The Advanced Modeling for Particle Accel...	DOE Mission Science	Provisioned
<input type="checkbox"/>	i	Physics : Cosmology	ERCAP0022176	2022	Dore, Olivier (dore)	NASA JPL	Simulating high-redshift line intensity ...	Cosmology	m3874	The project aims at simulating high reds...	DOE Mission Science	Provisioned
<input type="checkbox"/>	i	Physics : Cosmology	ERCAP0021824	2022	Nugent, Peter (nugent)	Berkeley Lab	DECaLS & the Zwicky Transient Facility	m937	m937	The Zwicky Transient Facility (next gene...	DOE Mission Science	Provisioned
<input type="checkbox"/>	i	Physics : High Energy Physics (Experimen...	ERCAP0020631	2022	Norman, Andrew (@norman)	Fermilab	Enabling HEP Intensity Frontier Science ...	FNAL-IF	m3249	The HEPCloud Facility is an intelligent ...	DOE Mission Science	Provisioned

(There's seldom a shortage of proposals from Berkeley/LBL)

And clicking on the shown project (chosen randomly)...

Rubin/DESC, as an example

ERCAP Number	ERCAP0020008	Allocation Year	2022
Project Title	LSST Dark Energy Science Collaboration (DESC) Image Simulations	State	Provisioned
Label	LSST_DESC	Project Class	DOE Mission Science ⓘ
PI Name	Digel, Seth (digel) ⓘ	Program	HEP - High Energy Physics ⓘ
ORG Name	SLAC ⓘ	Science Category	Physics : Cosmology ⓘ
Email	digel@slac.stanford.edu	Project	m1727
Business phone	650-926-3614		

Personnel	Funding	Security	Project Details	Resources	Codes	Supporting Information	Usage Agreement	Award Information
-----------	---------	----------	-----------------	-----------	-------	------------------------	-----------------	-------------------

Computational Resources - Enter the number of NERSC computational hours requested.

CPU Node Hours Used	5,537	CPU Node Hours Requested	13,000
GPU Node Hours Used	6,307	GPU Node Hours Requested	14,000

If you are applying for GPU time, please tell us about your [readiness](#).

GPU Readiness

Our main GPU-enabled code is the Batoid package developed by LSST DESC member Josh Meyers to do science for use with ImSim, <https://www.nersc.gov/research-and-development/hesaphesap-projects/#ImSim>, and he has demonstrated its readiness for use with ImSim. Our other GPU-enabled code, FlowPM, uses a standard TensorFlow framework and has already been run in production on clusters up to 256 GPUs.

A modest request for both CPU and GPU



Factors going into awarded allocations

- ▶ Here our guidelines
- ▶ We have obligations to Cosmic Frontier experiments which we must try to meet
 - ▶ because many have been urged over past years by HEP to use NERSC for their Production/Operations/Research computing to conserve resources
 - ▶ By agreement between HEP and ASCR ADs: Jim Siegrist and Barb Helland.
- ▶ We generally accommodate modest requests fully
 - ▶ Particularly if there is a history of using previous allocations
 - ▶ A lot of these are in the noise
- ▶ Historically large, efficient power users making large requests again are accommodated, as possible
 - ▶ Year-over-Year request over $\sim 30\%$ requires a conversation
- ▶ Large requests from users with a history of not using allocations in the past will frequently receive scaled back allocations in coming AY
- ▶ We run our proposal by the HEP Program Managers for concurrence
- ▶ We Hold Back $\sim 15\%$ for later allocation
- ▶ Similar consideration for GPU and storage
 - ▶ Knowing that GPU on Permuter was to be uncharged most of AY22



Ongoing adjustments

- ▶ During the year at a discrete point NERSC (Clayton Bagwell) makes their own adjustments, clawing back from under-users
- ▶ We come in around September and make our own adjustments
 - ▶ Clawing back from users under a certain CPU usage
 - ▶ Offering to add to those above some threshold
 - ▶ Points for maximizing charged-machine usage
 - ▶ Lots of communication back/forth before action
 - ▶ Good reasons for under-usage generally result in inaction
- ▶ We also do epsilon tuning to maximize usage in Nov/December
- ▶ This is done on the *iris* system



► Sorted by CPU node hours used

The lower this number the shrewder the use of queues and other discounts

← → ↻ iris.nersc.gov/user/96840/home

Search users, projects

Church, Eric

Project	Managers	Organization	Allocated	Charged	Raw Hours	Avg CF	Balance	% Used	CFS Storage
m3035	pi: Feng, Yu pi_proxy: Liu, Jia	UC Berkeley	500	519	2,790	0.93	-19	103.8%	150.0 TB
m3926	pi: Kronfeld, Andreas	Fermilab	140,000	143,588	1,423,161	0.50	-3,588	102.6%	484.0 GB
mp27	pi: Sinclair, Donald pi_proxy: Kogut, John	Argonne Lab	50,000	49,791	166,973	0.67	209	99.6%	3.2 TB
m3592	pi: Schlegel, David pi_proxy: Lang, Dustin	Berkeley Lab	110,000	103,321	430,380	0.44	6,679	93.9%	0.0 B
m4045	pi: Wilkinson, Callum	Berkeley Lab	80	72	360	1.00	8	90.0%	4.0 TB
deepsrch	pi: Perlmutter, Saul pi_proxy: Spadafora, Anthony pi_proxy: Huang, Xiaosheng	Berkeley Lab	1,000	860	2,476	0.98	140	86.0%	2.5 TB
m1647	pi_proxy: Gottlieb, Steven pi: El-Khadra, Aida pi_proxy: DeTar, Carleton pi_proxy: Mackenzie, Paul	U. Illinois U-C	169,000	125,070	1,086,268	0.58	43,930	74.0%	113.1 GB
atlas	pi_proxy: Hinchliffe, Ian pi: Marshall, Zachary pi_proxy: Wang, Haichen	Berkeley Lab	6,250	4,533	79,432	0.06	1,717	72.5%	565.0 TB
m349	pi: Ng, Cho-Kuen pi_proxy: Li, Zenghai	SLAC	2,500	1,753	5,682	0.99	747	70.1%	229.6 GB
m1779	pi: Ng, Cho-Kuen pi_proxy: Li, Zenghai	SLAC	5,000	3,387	11,678	1.00	1,613	67.7%	867.2 GB

Note that some users have pegged the needle already, 3mo's before AY end. Some groups' compute appetites are insatiable.

Storage increase requests: 20 Tbytes have been historically been granted mid-AY (not by us, by NERSC)

Out-Years NERSC usage

- ▶ With an eye toward gauging HEP demand
 - ▶ In the past year we have held discussions with ~ 8 of our historically heavy users, e.g.
 - ▶ LHC experiments
 - ▶ Lattice QCD projects
 - ▶ NERSC explicitly uninterested in our out-years
- ▶ We have asked these projects' anticipated usage for the next ~ 3 years
 - ▶ To the Point we've asked for this in their ERCAP proposals
 - ▶ And urged them toward ever-more GPU usage
 - ▶ Some groups are more prepared than others
 - ▶ A few have reported unreliable Perlmutter GPU running and expressed concern

